

ATTORNEY'S DOCKET NO. SGL-001
PATENTS

UNITED STATES PATENT APPLICATION

OF

SIMSON L. GARFINKEL

FOR

**SYSTEM AND METHOD THAT PROVIDES FOR THE EFFICIENT AND EFFECTIVE SANITIZING OF
DISK STORAGE UNITS AND THE LIKE**

Certificate of Express Mailing

Express Mail Mailing Label No. EK 904 503 507 US

Date of Deposit August 9, 2001

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office To Addressee" Service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D. C. 20231.

By *Richard Jordan*

Richard A. Jordan

1

FIELD OF THE INVENTION

2 The invention relates generally to the field of digital data processing and storage, and more
3 specifically to systems and methods that provide for sanitizing of disk storage units and the like.

4

BACKGROUND OF THE INVENTION

5 Digital computers, mass digital data storage subsystems and the like typically include disk
6 storage units to provide for relatively long-term storage of digital data. It is often necessary to
7 remove a disk storage unit from a computer, mass storage subsystem, etc., in which it formed
8 component. This may occur, for example, if it is necessary to remove the disk storage unit for repair,
9 if it is necessary to replace the disk storage unit, if the computer, mass storage subsystem, etc., is to
10 be discarded, or for other reasons that will be apparent to those skilled in the art.

11 Data stored in a disk storage unit is often confidential to the organization that maintains the
12 computer, mass storage subsystem, etc., in which the disk storage unit forms a component. Several
13 problems can arise in connection with maintaining the confidentiality of the data that is stored in a
14 disk storage unit. For example, disk storage units store data, organized into files, in magnetic form.
15 Typically, when a file is deleted, the data is not erased from the disk storage unit, but instead
16 information detailing the locations of the data comprising the respective file is deleted from tables
17 that are maintained therefor by the computer, mass storage subsystem, etc. Accordingly, merely

1 erasing files from a disk storage unit will not serve to erase the data contained in the files. The data
2 can be recovered using any of a number of conventional data recovery techniques.

3 Even if efforts are made to sanitize a disk storage unit, that is, to erase the data stored in a
4 disk storage unit, to over-write the data with other data, or to perform other sanitizing operations that
5 will be apparent to those skilled in the art, it is often still possible to recover the erased or over-
6 written data, since it is not unusual for residual magnetic fields to remain after the erasure or over-
7 writing that can be detected sufficiently for the data represented thereby to be reconstructed.

8 In addition, since the data storage capacity of disk storage units is quite large and growing, the time
9 required to over-write the data stored on a typical disk storage unit even once is prohibitive, and
10 typically data is not considered "wiped" until it has been over-written at least several times, generally
11 with predetermined data patterns.
12
13
14

15 Moreover, during wiping, a "Trojan horse" program can cause data to be copied from the
16 storage locations in which it is currently stored to spare storage locations on the disk storage unit that
17 may be provided to accommodate the possibility that some of the "regular" storage locations may
18 go bad. If a regular storage location does go bad, the disk storage unit automatically stores the data
19 that is to be stored on the bad regular storage location on a spare storage location that has been
20 allocated therefor. Thereafter, when the data is to be retrieved from a "regular" storage location for
21 which a spare storage location has been allocated, the disk storage unit will automatically retrieve
the data from the spare storage location and provide the data to the device that requested the data.
Generally, the spare storage locations will be known to the disk storage unit, and not to the device,
that is, the computer or the like that stores data in, and retrieves data from the disk storage unit, and

1 so the wiping will be in connection with the regular storage locations and not the spare storage
2 locations. In that case, the data will still be available in the spare storage locations.

3 Instead of overwriting or wiping a disk storage unit, the contents of a disk storage unit can
4 be erased in a "bulk erasure" operation by bringing the disk storage unit in close proximity to a
5 strong magnetic field to "de-gauss" the disk storage unit. However, de-gaussing a disk storage unit,
6 in addition to erasing the data stored thereon, will also erase formatting information that identifies
7 the storage locations, making the disk storage unit thereafter unusable.

8
9
10
11
12
13
14
15
16
17
18
19

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for the efficient and effective sanitizing of disk storage units, while additionally providing that the disk storage units can be subsequently used without risking the confidentiality of data previously stored thereon.

In brief summary, generally the invention provides for sanitizing of a digital data storage unit, such as a disk data storage unit, by encrypting the information that is stored thereon. When the previously-stored information is retrieved from the digital data storage system storage unit, the information is decrypted prior to being provided to the device that requested retrieval of the data. If the digital data storage unit is to be sanitized, the key or keys that used to at least decrypt the data stored on the digital data storage unit are discarded or made unavailable for use in decrypting the encrypted digital data, thereby making the unencrypted data unavailable. Any of a number of types of encryption/decryption methodologies can be used, including a symmetric key methodology, an

1 asymmetric key methodology such as a public key/private key methodology, or any of a number of
2 other encryption/decryption methodologies as will be apparent to those skilled in the art.

3 The decryption key can be provided to the disk storage unit in any of a number of ways that
4 would facilitate discarding of the decryption key when the disk storage unit is to be sanitized. For
5 example, the decryption key may be stored in non-volatile random-access memory ("NVRAM") and
6 available for use in decrypting data retrieved from the disk storage unit. When the disk storage unit
7 is to be sanitized, the NVRAM can be made unavailable to the person or entity that has possession
8 of the disk storage unit, by, for example, discarding the NVRAM, erasing the NVRAM, destroying
9 the NVRAM, or any other arrangement whereby the contents of the NVRAM are unavailable to the
10 person or entity that has possession of the disk storage unit. Since this makes the decryption key
11 unavailable to the person or entity that has possession of the disk storage unit, the person or entity
12 that has possession of the disk storage unit will be unable to retrieve the data in unencrypted form.
13
14

15 Alternatively, some or all of the decryption key can be stored on the disk storage unit itself,
16 and another portion provided in, for example, an integrated circuit ("IC") chip, and the key obtained
17 by performing a selected processing operation in connection with the portion of the key stored on
18 the disk storage unit and the portion stored on the IC chip. For example, the decryption key can be
19 formed by concatenating the portion stored on the IC chip with the portion stored on the disk storage
20 unit. Alternatively, the decryption key can be formed by performing an exclusive-OR ("XOR")
21 operation in connection with the portion stored on the disk storage unit and the portion stored on the
IC chip. Other types of processing operations that can be used in connection with the portion stored
on the disk storage unit and the portion stored on the IC chip, which can be used to generate a

1 decryption key, will be apparent to those skilled in the art. When the disk storage unit is to be
2 sanitized, the IC chip can be removed, destroyed, or the like, which can make it impossible to
3 reconstruct the key(s).

4 **BRIEF DESCRIPTION OF THE DRAWINGS**

5 This invention is pointed out with particularity in the appended claims. The above and
6 further advantages of this invention may be better understood by referring to the following
7 description taken in conjunction with the accompanying drawings, in which:

8 FIG. 1 is a functional block diagram of a mass storage subsystem including a system that
9 provides for the efficient and effective sanitizing of disk storage units included therein, constructed
10 in accordance with the invention; and

11 FIG. 2 is a functional block diagram of a second embodiment of a mass storage subsystem
12 including a system that provides for the efficient and effective sanitizing of disk storage units
13 included therein, constructed in accordance with the invention.

14 **DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT**

15 FIG. 1 is a functional block diagram of a mass storage subsystem 10 including a system that
16 provides for the efficient and effective sanitizing of disk storage units included therein, constructed
17 in accordance with the invention. With reference to FIG. 1, the mass storage subsystem 10 includes

1 one or more disk storage units 11(1) through 11(N) (generally identified by reference numeral
2 11(n)), a controller 12, an interface 13 and one or more key stores 14(1) through 14(N) (generally
3 identified by reference numeral 14(n)). Each disk storage unit 11(n) stores digital data provided to
4 the mass storage subsystem 10 by an external data utilization device (not shown), which digital data
5 can later be retrieved, by the same or a different data utilization device, for usage. The data
6 utilization device(s) may include any of a number of types of devices, including, but not limited to,
7 computers, including personal computers, computer workstations, mini- and mainframe computers,
8 that retrieve data from the mass storage subsystem 10, process the data, and provide processed data
9 to the mass storage subsystem 10 for subsequent storage. In addition, the data utilization device(s)
10 may include any of a number of types of other devices, including data visualization devices that
11 display the data, in text, image or any other form, to a user, devices for producing hardcopy output,
12 network or other communication devices for transferring data over a network or other
13 communication medium to other data utilization devices, backup devices for producing backup
14 copies of data stored on the mass storage subsystem 10, as well as any of a number of other types
15 of devices that can produce and/or make use of digital data.

16 The interface 13 connects to one or more of the data utilization devices over a
17 communication link 16. The communication link 16 may be any of a number of types of
18 communication links over which information in digital form may be transferred. The interface 13
19 receives storage retrieval requests from data utilization devices over the communication link 16. A
20 storage request initiates a storage operation in connection with the mass storage subsystem 10, to
21 facilitate the storage of data on one or more of the disk storage units 11(n). Storage requests may

1 be accompanied by the data to be stored, or the data may be provided to the mass storage subsystem
2 10 separately from the request. A retrieval request initiates a retrieval operation to enable data to
3 be retrieved from one or more of the disk storage units 11(n). After the data has been retrieved, it
4 can be transferred by the interface 13 over the communication link 16 to the device that issued the
5 retrieval request, or to one or more other devices as directed by the retrieval request.

6 After the interface 13 has received a storage or retrieval request, it will provide the request
7 to the controller 12 for processing. The controller performs a number of operations in connection
8 with the mass storage subsystem 12, including scheduling of storage and retrieval operations by the
9 respective disk storage units 11(n), buffering of data to be stored in a storage operation pending
10 storage in a disk storage unit 11(n), buffering of data retrieved from a disk storage unit 11(n) prior
11 to transmission to the destination data utilization device by the interface 13, and the like.

12 In addition, the controller 12 includes an encryption/decryption module 15 that performs an
13 encryption operation in connection with data to be stored on a disk storage unit 11(n) to encrypt the
14 data before it is transferred thereto for storage, and a decryption output in connection with data
15 retrieved from a disk storage unit 11(n) to decrypt the data after it has been retrieved and before it
16 is transferred by the interface 13 to the destination data utilization device. When performing an
17 encryption or decryption operate in connection with data to be stored on or that has been retrieved
18 from a disk storage unit 11(n), the encryption/decryption module 15 makes use of a key stored in a
19 correspondingly-indexed key store 14(n). The encryption/decryption module 15 may make use of
20 any of a number of encryption/decryption methodologies, including a symmetric key methodology,
21 an asymmetric key methodology such as a public key/private key methodology, or any of a number

1 of other encryption/decryption methodologies as will be apparent to those skilled in the art. If the
2 encryption/decryption module 15 makes use of a symmetric key methodology in connection with
3 a disk storage unit 11(n), the key store 14(n) will store one key that it will use for both encryption
4 of data to be stored and decryption of data that has been retrieved. On the other hand, if the
5 encryption/decryption module makes use of an asymmetric key methodology in connection with a
6 disk storage unit 11(n), the key store 14(n) will store two keys, namely, an encryption key that it will
7 use for encryption of data to be stored and a decryption key that it will use for decryption of data that
8 has been retrieved. It will be appreciated that the encryption/decryption module 15 may make use
9 of different methodologies for different ones of the disk storage units 11(n), in which case one or
10 more of the key stores 14(n) may store one key used for both encryption and decryption, and others
11 may store two keys, one used for encryption and the other used for decryption. In addition, it will
12 be appreciated that data stored on one or more of the disk storage units 11(n) may not be encrypted,
13 in which case no key need be provided therefor.

14 As noted above, one or more of the disk storage units 11(n) generally store data provided
15 thereto by the controller 12 in encrypted form. For the disk storage units 11(n) for which data is
16 stored in encrypted form, as long as the key is available to decrypt the data stored on the disk storage
17 unit 11(n), the data stored on the disk storage unit 11(n) can be decrypted and provided to a data
18 utilization device. Accordingly, if it is desired to make the data stored on a disk storage unit 11(n)
19 unavailable for access, at least in unencrypted form, the key that is used for decryption is discarded.
20 This can be accomplished in a number of ways. For example, the key store 14(n) can merely be
21 erased, if it the key store 14(n) is in the form of a volatile memory. Alternatively, if the key store

1 14(n) is in the form of a non-volatile memory, the key store 14(n) can be removed from the mass
2 storage subsystem 10. In that case, erasure or destruction of the non-volatile memory comprising
3 the key store 14(n) can ensure that the data stored on the disk storage unit 11(n) will remain
4 unavailable for access in unencrypted form; the non-volatile memory can be erased using any of a
5 number of conventional erasure techniques, including, for example, application of a selected voltage
6 to the circuitry comprising the key store. As another alternative, the key store 14(n) can comprise
7 a selected storage location on the disk storage unit 11(n) itself, in which case the key can be erased
8 by erasing the selected storage location. As yet another alternative, the key or keys to be used in
9 connection with a disk storage unit 11(n) can initially be stored on the disk storage unit, and
10 thereafter copied by the controller 12 to a separate volatile or non-volatile key store 14(n); in that
11 case, when it is desired to make the data on the disk storage unit 11(n) unavailable in unencrypted
12 form, operations can be performed in a manner similar to those described above in connection with
13 both the disk storage unit 11(n) and the key store 14(n).

14 If the key or keys used by the encryption/decryption module 15 to encrypt data to be stored
15 on a disk storage unit 11(n), and to decrypt data retrieved from the respective disk storage unit 11(n),
16 are stored in a key store 14(n) separate and apart from the respective disk storage unit 11(n), the key
17 or keys may be initially provided in a number of ways. For example, the key(s) can be initially
18 stored on the disk storage unit 11(n) and copied by the controller 12 to the respective key store 14(n)
19 as part of an initialization procedure when the disk storage unit is first installed in the mass storage
20 subsystem 10, when the mass storage subsystem 10 is first powered on or configured, or as part of
21 any other procedure as will be apparent to those skilled in the art. The controller 12 may also enable

1 the disk storage unit 11(n) to erase the key(s) from the disk storage unit 11(n) if, for example, the
2 key store 14(n) is a non-volatile memory. Alternatively or in addition, the key(s) may be provided
3 in machine readable form on another machine readable medium that may be read by the controller
4 12 using a suitable reading device (not separately shown).

5 As a further alternative, the controller may be provided with a key generator module that can
6 generate the key(s) to be used with the disk storage units. One illustrative embodiment thereof will
7 be described in connection with FIG. 2. FIG. 2 depicts a mass storage subsystem 110 that includes
8 one or more disk storage units 111(n), an interface 113 connected to a communication link 116, key
9 store(s) 114(n) and a controller 112 that generally correspond to and operate in a manner similar to
10 respective disk storage unit(s) 11(n), interface 13, communication link 116, key stores 14(n) and
11 controller 112, except as follows.

12 In addition to an encryption/decryption module 115, which corresponds to and operates in
13 a manner similar to encryption/decryption module 15, the controller 112 is also provided with a key
14 generator 117 that generates one or more keys for a disk storage unit 111(n) by use of a bit pattern
15 stored on, for example, the respective disk storage unit 111(n) itself and another bit pattern. The bit
16 pattern stored on the respective disk storage unit 111(n) may be stored on, for example, a storage
17 location on the disk storage unit identified as pattern store 119(n). The other bit pattern used by the
18 key generator 117 in generating one or more keys for the disk storage unit 111(n) may be provided
19 in another pattern store 118(n), which may comprise a component of the mass storage subsystem
20 110, a machine-readable medium that is readable by an appropriate reading device, or the like. The
21 mass storage subsystem 110 may include a single pattern store for storing a bit pattern that can be

1 used for all disk storage unit(s) included therein, a plurality of pattern stores 118(n) each of which
2 stores a bit pattern for a correspondingly-indexed disk storage unit 111(n), or a plurality of pattern
3 stores each of which stores a bit pattern that will be used in connection with generating one or more
4 keys for selected ones of the disk storage units 111(n). The key generator 117 can perform any of
5 a number of processing operations in connection with the bit patterns, which may include
6 concatenating the bit patterns together, exclusive-ORing ("XORing") some or all of the bit patterns
7 together, or any other types of processing operations as will be apparent to those skilled in the art.

8 When it is desired to sanitize a disk storage unit 111(n) in the mass storage subsystem 110,
9 the decryption key can be made unavailable by erasing or destroying the key store 114(n) and erasing
10 or destroying one or both of the pattern stores 118(n), 119(n), thereby to make one or both of the bit
11 patterns stored therein unavailable. Making one of the bit patterns unavailable will generally suffice
12 to make regeneration of the key(s) difficult if not impossible.

13 The invention provides a number of advantages. In particular, the invention provides an
14 arrangement whereby a disk storage unit can be efficiently and effectively sanitized in such a manner
15 that it can subsequently provided to a repair facility, be used in another environment, and the like,
16 while maintaining the confidentiality of the data stored thereon. This is accomplished by providing
17 that the data stored on the disk storage unit is encrypted, and decrypted during a retrieval operation,
18 and further providing that the key used in decrypting the data be unavailable when it is desired to
19 sanitize the disk storage unit. The key used in decrypting the data can be made unavailable by, for
20 example, erasing or destroying a store in which the key is stored; if the key is stored in multiple
21 stores, preferably all of the stores will be, for example, erased or destroyed so that the key will

1 thereafter be unavailable for use in decryption. Without the key, the data stored on the disk storage
2 unit generally cannot be retrieved in unencrypted form, in which case the data will remain
3 confidential if the disk storage unit is provided to an entity other than that for which the data
4 comprises confidential information.

5 It will be appreciated that a number of changes and modifications may be made to the
6 arrangement described herein. As noted above, any of a number of conventional
7 encryption/decryption methodologies may be used, including both symmetric key and asymmetric
8 key methodologies. It will be appreciated that generally the ability of an encryption/decryption
9 methodology to maintain the confidentiality of the encrypted data will reflect the security of the
10 methodology and the number of bits comprising the key(s) used in encryption and decryption. In
11 addition, it will be appreciated that, if it is desired to have the data stored on the disk storage unit
12 available in unencrypted form after the disk storage unit has been sanitized, it will generally be
13 desirable to backup the data before the disk storage unit is sanitized, using any of a number of
14 conventional backup arrangements. The data may be backed up directly onto another disk storage
15 unit, or onto a backup medium such as tape, after which it can be loaded onto another disk storage
16 unit for subsequent retrieval.

17 Although the arrangement has been described as being used in connection with disk storage
18 units, it will be appreciated that the arrangement can also be used in connection with other types of
19 digital data storage arrangements, including storage subsystems that emulate disk storage units but
20 make use of storage media other than magnetic disks.

1 It will be appreciated that a system in accordance with the invention can be constructed in
2 whole or in part from special purpose hardware or a general purpose computer system, or any
3 combination thereof, any portion of which may be controlled by a suitable program. Any program
4 may in whole or in part comprise part of or be stored on the system in a conventional manner, or it
5 may in whole or in part be provided in to the system over a network or other mechanism for
6 transferring information in a conventional manner. In addition, it will be appreciated that the system
7 may be operated and/or otherwise controlled by means of information provided by an operator using
8 operator input elements (not shown) which may be connected directly to the system or which may
9 transfer the information to the system over a network or other mechanism for transferring
10 information in a conventional manner.

11 The foregoing description has been limited to a specific embodiment of this invention. It will
12 be apparent, however, that various variations and modifications may be made to the invention, with
13 the attainment of some or all of the advantages of the invention. It is the object of the appended
14 claims to cover these and such other variations and modifications as come within the true spirit and
15 scope of the invention.

16 What is claimed as new and desired to be secured by Letters Patent of the United States is: